Cognitive Illusions of Authorship Reveal Hierarchical Error Detection in Skilled Typists

Gordon D. Logan, *et al.*
Science **330**, 683 (2010);
DOI: 10.1126/science.1190483

ited by a daughter progressively amplifies to a steady-state plateau as that cell ages (Fig. 3E). Thus, early daughters will inherit fewer aggregates than later daughters. Indeed, propagons transmitted to successive daughters increased through the third generation and then remained relatively constant (Fig. 3F). Our observations provide a mechanistic explanation for previously observed cell-to-cell variability in propagons (*6*, *25*) and reveal age-dependent [*PSI*$^+$] phenotypes at the single-cell level.

Thus, the seemingly static phenotypes associated with prion protein conformations may actually reflect highly dynamic pathways of prion protein biogenesis in dividing cells. For any given cell, the complement of aggregates and the phenotype fluctuate in response to the interplay between the protein-misfolding pathway and its cellular environment, creating a self-regenerating system that settles to a stable population average for each conformation. Thus, the cellular environment has profound effects on the phenotypic manifestations of prion protein conformations.

The dynamic size-based system that we have uncovered may contribute to the physiological consequences of protein misfolding in ways that are not possible for an abundance-based process. The phenotypic variation established and maintained through the events described here strengthens the argument that the prion mechanism, like other epigenetic processes, facilitates selection in new environments and consequently evolution (*26*). According to our model, access to an advantageous state may not require a [*prion*$^-$]/[*PRION*$^+$] phenotypic switch (*27*) but instead may always be present within a population. In mammals, variation in aggregate size may similarly affect protein transmissibility between nondividing cells and the spread of pathology in prion and perhaps other protein-misfolding disorders, such as Parkinson's, Alzheimer's, and Huntington's diseases (*28*–*30*). Indeed, prion protein conformation and expression—parameters that alter aggregate size—are more reliable predictors of the clinical course of disease than is the presence of protease-resistant aggregates (*1*).

References and Notes

1. J. Collinge, A. R. Clarke, *Science* **318**, 930 (2007).
2. L. Benkemoun, S. J. Saupe, *Fungal Genet. Biol.* **43**, 789 (2006).
3. S. B. Prusiner, *Science* **216**, 136 (1982).
4. J. A. Pezza, T. R. Serio, *Prion* **1**, 36 (2007).
5. M. Tanaka, S. R. Collins, B. H. Toyama, J. S. Weissman, *Nature* **442**, 585 (2006).
6. B. Cox, F. Ness, M. Tuite, *Genetics* **165**, 23 (2003).
7. S. S. Sindi, T. R. Serio, *Curr. Opin. Microbiol.* **12**, 623 (2009).
8. I. L. Derkatch, Y. O. Chernoff, V. V. Kushnirov, S. G. Inge-Vechtomov, S. W. Liebman, *Genetics* **144**, 1375 (1996).
9. S. M. Uptain, G. J. Sawicki, B. Caughey, S. Lindquist, *EMBO J.* **20**, 6236 (2001).
10. P. Zhou *et al.*, *EMBO J.* **18**, 1182 (1999).
11. D. S. Kryndushkin, I. M. Alexandrov, M. D. Ter-Avanesyan, V. V. Kushnirov, *J. Biol. Chem.* **278**, 49636 (2003).
12. S. N. Bagriantsev, V. V. Kushnirov, S. W. Liebman, *Methods Enzymol.* **412**, 33 (2006).
13. Materials and methods are available as supporting material on *Science* Online.
14. F. Ness, P. Ferreira, B. S. Cox, M. F. Tuite, *Mol. Cell. Biol.* **22**, 5593 (2002).
15. P. Satpute-Krishnan, S. X. Langseth, T. R. Serio, *PLoS Biol.* **5**, e24 (2007).
16. T. Haslberger, B. Bukau, A. Mogk, *Biochem. Cell Biol.* **88**, 63 (2010).
17. H. K. Edskes, V. T. Gray, R. B. Wickner, *Proc. Natl. Acad. Sci. U.S.A.* **96**, 1498 (1999).
18. K. D. Allen *et al.*, *Genetics* **169**, 1227 (2005).
19. M. Crapeau, C. Marchal, C. Cullin, L. Maillet, *Mol. Biol. Cell* **20**, 2286 (2009).
20. A. B. Salnikova, D. S. Kryndushkin, V. N. Smirnov, V. V. Kushnirov, M. D. Ter-Avanesyan, *J. Biol. Chem.* **280**, 8808 (2005).
21. E. G. Rikhvanov, N. V. Romanova, Y. O. Chernoff, *Prion* **1**, 217 (2007).
22. V. V. Kushnirov, A. B. Vishnevskaya, I. M. Alexandrov, M. D. Ter-Avanesyan, *Prion* **1**, 179 (2007).
23. T. R. Serio, S. L. Lindquist, *Annu. Rev. Cell Dev. Biol.* **15**, 661 (1999).
24. B. Cox, *Heredity* **20**, 505 (1965).
25. L. J. Byrne *et al.*, *PLoS ONE* **4**, e4670 (2009).
26. R. Halfmann, S. Alberti, S. Lindquist, *Trends Cell Biol.* **20**, 125 (2010).
27. A. K. Lancaster, J. P. Bardill, H. L. True, J. Masel, *Genetics* **184**, 393 (2010).
28. J. R. Silveira *et al.*, *Nature* **437**, 257 (2005).
29. B. Frost, M. I. Diamond, *Nat. Rev. Neurosci.* **11**, 155 (2010).
30. P. Brundin, R. Melki, R. Kopito, *Nat. Rev. Mol. Cell Biol.* **11**, 301 (2010).
31. J. A. Pezza *et al.*, *Mol. Biol. Cell* **20**, 1068 (2009).

32. We thank N. Rhind, M. Duennwald, J. Sherley, and members of their groups for assistance with centrifugal elutriation, and the Center for Computation and Visualization, the Leduc Bioimaging Facility, and the Flow Cytometry and Sorting Facility at Brown University for instrumentation and the assistance of the respective support staffs. We thank J. Bender, J. Laney, M. Tuite, and members of the Serio and Laney labs for helpful discussions and comments on the manuscript. This research was supported by grants from NIH (GM069802 to T.R.S., GM085976 to A.D., GM089049 to S.S.S., AG034754 to C.K., and AG032818 to S.D.) and NSF/ADVANCE (0548311 to T.R.S.).

Supporting Online Material

www.sciencemag.org/cgi/content/full/330/6004/680/DC1
Materials and Methods
Figs. S1 to S12
Tables S1 and S2
References

14 September 2010; accepted 5 October 2010
10.1126/science.1197785

Cognitive Illusions of Authorship Reveal Hierarchical Error Detection in Skilled Typists

Gordon D. Logan* and Matthew J. C. Crump

The ability to detect errors is an essential component of cognitive control. Studies of error detection in humans typically use simple tasks and propose single-process theories of detection. We examined error detection by skilled typists and found illusions of authorship that provide evidence for two error-detection processes. We corrected errors that typists made and inserted errors in correct responses. When asked to report errors, typists took credit for corrected errors and accepted blame for inserted errors, claiming authorship for the appearance of the screen. However, their typing rate showed no evidence of these illusions, slowing down after corrected errors but not after inserted errors. This dissociation suggests two error-detection processes: one sensitive to the appearance of the screen and the other sensitive to keystrokes.

Errors are ubiquitous in human performance (*1*, *2*). Their consequences can be costly, ranging from mild annoyance to global-scale disaster. Errors are common in the early stages of skill acquisition, when learning is based on trial and error, but they still prevail in expert performance, when skills are automatic and fluent. Detecting and correcting errors are important components of executive control at high skill levels (*3*, *4*). The control processes that manage errors are evident in behavioral measures of post-error slowing and conscious reports of errors (*1*), as well as in neural measures of error-related potentials in the electroencephalogram (*5*) and activation of the anterior cingulate cortex in functional brain imaging (*6*, *7*). Several theories have been proposed to account for these measures. Some assume a post hoc comparison of intended and actual actions (*5*). Others suggest that conflict between competing responses may be sufficient to trigger error detection and correction (*4*). Researchers often view these accounts as mutually exclusive and perform experiments intended to decide between them, as if there were a single error-detection mechanism and the goal of their research was to determine its properties. That may be possible in simple tasks with single responses, but it is unlikely to apply to complex tasks, like typewriting, that involve hierarchical processing and extended interaction with the environment (*8*–*10*). Here, we report the induction of cognitive illusions of authorship (*11*–*14*) in skilled typists and provide evidence for two separate error-detection mechanisms nested in a hierarchical control process (*10*).

Skilled typewriting engages many processes, from perception to cognition and action (*9*, *15*). The processes can be divided into two nested

Department of Psychology, Vanderbilt University, Nashville, TN 37203, USA.

*To whom correspondence should be addressed. E-mail: gordon.logan@vanderbilt.edu

feedback loops: (i) an outer loop that begins with language comprehension and generation and produces a series of words to be typed and (ii) an inner loop that begins with the words the outer loop provides and produces a series of keystrokes (*9*, *10*, *15*–*17*). The two loops are relatively autonomous: The outer loop commands the inner loop but knows little about how the inner loop carries out its commands (*10*, *17*). Our research was guided by the hypothesis that the two loops are driven by different kinds of feedback: The outer loop is sensitive to the appearance of the output on the screen, whereas the inner loop is sensitive to feedback from the fingers and the keyboard (*9*, *17*, *18*). Thus, the two loops should detect errors with different mechanisms (*19*, *20*).

We created illusions of authorship by introducing mismatches between errors in the inner and outer loops. We corrected errors that typists made, so the output matched their intentions when their motor behavior did not, and we inserted errors into the output that the typists produced, so that their motor behavior matched their intentions when the output did not. We hypothesized that the outer loop would evaluate the match between intended and actual actions, claiming authorship for the appearance of the screen. The outer loop decides that the response is correct if the screen looks right and incorrect if the screen looks wrong. Thus, the outer loop would treat corrected errors like actual, typist-produced correct responses and inserted errors like actual, typist-produced errors, creating illusions of authorship (*11*–*14*). We hypothesized that the inner loop would monitor keystrokes, evaluating proprioceptive and kinesthetic feedback (*17*–*20*), and so would respond differently to inserted errors and actual errors and to corrected errors and correct responses. The inner loop would know the truth behind the illusion.

We tested these hypotheses in three experiments in which skilled typists typed single words presented one at a time on a computer screen, with their responses echoed on the screen below the word to be typed (*21*). We measured outer-loop error detection by asking for explicit reports of errors. The three experiments varied in the intrusiveness of explicit error detection. The first was the least intrusive, asking for retrospective reports after the experiment finished. The second asked participants to judge whether each word was correct or erroneous as soon as they typed it. The third told participants about corrected and inserted errors at the outset and asked them to judge whether each word was correct, erroneous, a corrected error, or an inserted error. We measured inner-loop error detection by evaluating post-error slowing. Skilled typists typically slow down immediately after an error, prolonging the interkeystroke interval between the error and the next keystroke (*18*). We predicted post-error slowing after actual errors and corrected errors and no slowing for inserted errors and correct responses.

We assessed the illusion of authorship by comparing explicit error detection with post-error slowing. If outer- and inner-loop error detection are accomplished by a single process, then typists should report errors in conditions that produce post-error slowing (actual errors and corrected errors), and they should report correct responses in conditions that do not produce post-error slowing (correct responses and inserted errors). If outer- and inner-loop error detection involve separate, hierarchically nested processes, then typists should report "correct" whenever the screen looks right (correct responses and corrected errors) and "error" whenever the screen looks wrong (actual errors and inserted errors), whether or not those conditions produce post-error slowing.

In each experiment, we inserted errors on 6% of the trials. On the remaining 94% of trials, we corrected ~45% of actual errors by echoing the correct response on the computer screen regard-

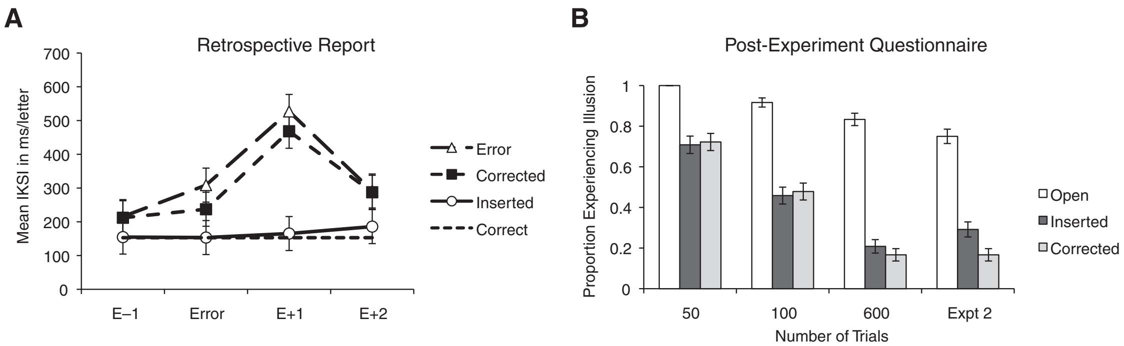

Fig. 1. (**A**) Mean interkeystroke interval in milliseconds per letter for the trial preceding an error (E–1), the error trial (Error), and the two trials after the error (E+1, E+2) for the 600-word group in experiment one, which used retrospective error reports gathered in a questionnaire administered after the experiment. Error bars are the 95% confidence intervals for the means based on Fisher's least significant difference test, calculated from the interaction between response type (error, inserted, corrected) and error position (E–1, Error, E+1, E+2). The data for correct responses are averaged over positions. (**B**) Proportions of participants who experienced illusions of authorship (did not report noticing inserted and corrected errors in question two, which was open-ended, in question three, which addressed inserted errors, and in question four, which addressed corrected errors) in the 50-, 100-, and 600-word groups and in experiment two. Error bars are 95% confidence intervals for the proportions calculated from the binomial distribution.

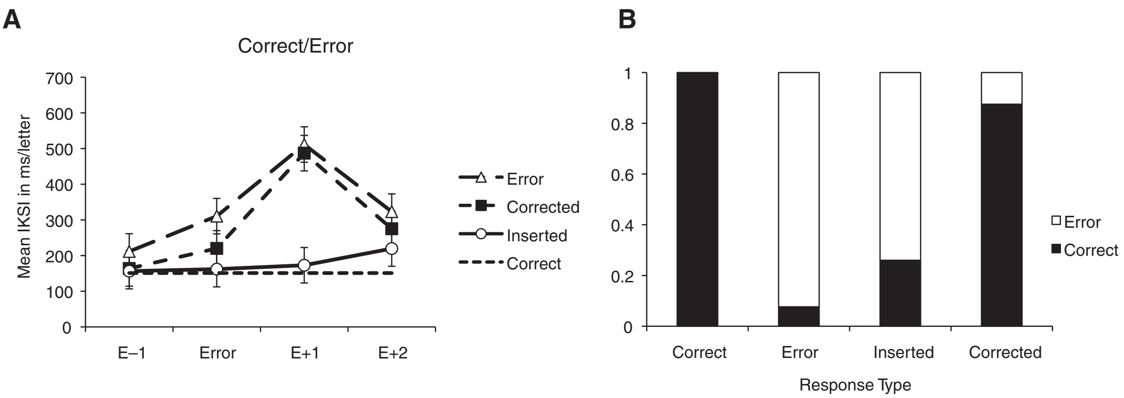

Fig. 2. (**A**) Mean interkeystroke interval in milliseconds per letter for the trial preceding an error (E–1), the error trial (Error), and the two trials after the error (E+1, E+2) in the second experiment, which used post-trial error reports in which participants said whether the preceding response was correct or incorrect. The data for correct responses are averaged over positions. Error bars indicate 95% confidence intervals based on Fisher's least significant difference test calculated from the interaction between response type and error position. (**B**) Mean probabilities of reporting correct or error for correct responses (correct), actual errors (error), inserted errors (inserted), and corrected errors (corrected).

less of what participants typed on 45% of the trials. We quantified skill as words per minute (WPM) on a typing test administered at the end of the experiment, reporting mean WPM ± standard deviation for each experiment. The typists were predominantly of college age, had 12.86 ± 5.35 years of experience, and typed at speeds comparable to professional typists.

The first experiment used retrospective error-detection reports, varying the number of words that were typed (50, 100, or 600) between participants to manipulate the opportunity to detect corrected and inserted errors. We tested three different groups of 24 skilled typists (WPM = 61.6 ± 16.9, 70.6 ± 20.6, and 68.6 ± 20.7 words, respectively); they typed 91.0, 89.9, and 91.0% of the words correctly, respectively. We assessed post-error slowing by examining interkeystroke interval for the trial before and two trials after an error. Twenty-two typists in the 600-word group provided sufficient data for this analysis. Their results, plotted in Fig. 1A, show slowing immediately after the error for incorrect responses ($F_{1,126}$ = 150.8, p < 0.01) and corrected errors ($F_{1,126}$ = 114.8, p < 0.01), but no slowing for inserted errors (F < 1.0). Thus, the inner loop responds differently to actual and inserted errors, as predicted. Interkeystroke intervals were longer for corrected errors and actual errors than for inserted errors and correct responses. This difference may reflect early error detection in the inner loop (*20*) or differences between words typed correctly and incorrectly. Errors are more likely to occur in words that are more difficult to type, and so are typed more slowly (*21*).

We assessed illusions of authorship with a post-experiment questionnaire, which consisted of six questions. The second, third, and fourth questions were most relevant. The first question asked for an estimate of the number of correct and incorrect responses. The second was open-ended, asking, "Did you notice anything about the kind of errors that you made?" Responses to this question were scored as illusions of authorship if typists made no mention of inserted or corrected errors. The third and fourth questions were direct, asking, "Did you notice that on some proportion of the trials the computer may have inserted errors, even if you correctly typed the word?" and "Did you notice that on some proportion of the trials the computer may have correctly typed a word even though you made an error?" Responses to these questions were scored as illusions of authorship if typists said "no."

The proportions of typists showing illusions of authorship are plotted in Fig. 1B. The proportions were significantly greater than zero for each question in each group (that is, the 95% confidence intervals constructed from the binomial distribution did not include 0). Illusions of authorship were stronger in the open-ended question than in the more direct questions, perhaps because typists were less willing to admit they had not noticed inserted and corrected errors when they were told they had occurred. Illusions of authorship were about the same for inserted and corrected errors. Illusions of authorship declined as typists experienced more inserted and corrected errors. There were 3, 6, and 36 inserted errors in the 50-, 100-, and 600-word versions of the experiment, respectively. Corrected errors were generated by randomly correcting 45% of all responses. The average numbers of corrected errors for the 18, 23, and 24 typists who experienced one or more in the 50-, 100-, and 600-word versions of the experiment were 2.4, 4.3, and 24, respectively.

The post-error–slowing data and retrospective reports show a dissociation between inner- and outer-loop error detection. Typists slowed after corrected errors but not after inserted errors, yet many of them accepted corrected errors as correct responses and inserted errors as errors. This analysis of post-error slowing collapses over typists who did and did not show illusions of authorship. We analyzed post-error slowing separately for typists who did and did not show illusions and found the same pattern in both groups (fig. S1) (*21*). Thus, the pattern in Fig. 1A is representative of all typists.

The second experiment asked 24 skilled typists (WPM = 68.2 ± 11.0) to type 600 words. After each word, they were asked to report errors explicitly, indicating whether they typed the word correctly or incorrectly. There was no mention of corrected and inserted errors. Typists typed 88.4% of the words correctly. Mean interkeystroke intervals from 23 typists who provided sufficient data for analysis are plotted in Fig. 2A. There was post-error slowing for incorrect responses ($F_{1,132}$ = 125.0, p < 0.01) and corrected errors ($F_{1,132}$ = 169.6, p < 0.01), but not for inserted errors (F < 1.0), suggesting again that inner-loop error detection distinguishes between actual errors and correct responses.

Explicit error-detection probabilities, plotted in Fig. 2B, show good discrimination between correct and incorrect responses. For correct responses, typists said "correct" more than "error" [$t(23)$ = 41.61, p < 0.01]; for incorrect responses, typists said "error" more than "correct" [$t(23)$ = 10.40, p < 0.01]. Typists showed illusions of authorship for inserted errors, saying "error" more than "correct" [$t(23)$ = 5.33, p < 0.01]. They also showed illusions of authorship for corrected errors, saying "correct" more than "error" [$t(23)$ = 5.26, p < 0.01].

The post-error slowing and post-trial report data reveal a dissociation between inner- and outer-loop error detection. We assessed the dissociation further by comparing trials in which typists did and did not experience illusions of authorship (*21*). The pattern of post-error slowing was the same for both sets of trials (fig. S5), suggesting that the pattern in Fig. 2A is representative of all trials.

The explicit report task did not allow typists to distinguish between actual errors and inserted errors or between correct responses and corrected errors. Nevertheless, typists chose responses that reflected the appearance of the screen instead of the keys they struck, consistent with their experience with computers and with our hypothesis that explicit error detection reflects outer-loop control processes. We administered the post-experiment questionnaire from experiment one and found that typists did not seem to be confused about the requirement to classify inserted and corrected errors. The proportions of typists showing illusions of authorship were about the same as those of the typists in the 600-word group in experiment one, who were not required to report errors after each trial (Fig. 1B).

We conducted a third experiment that included four response categories in the explicit detection

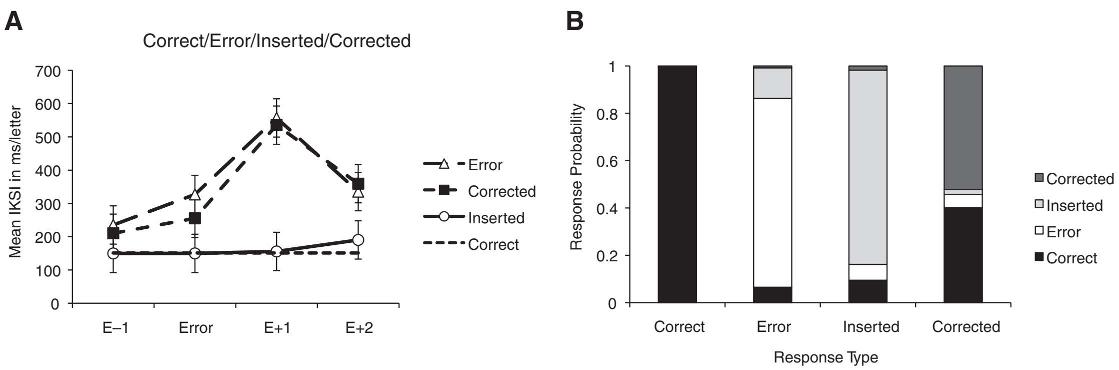

Fig. 3. (**A**) Mean interkeystroke interval in milliseconds per letter for the trial preceding an error (E–1), the error trial (Error), and the two trials after the error (E+1, E+2) in the third experiment, which used post-trial error reports in which participants said whether the preceding response was correct, error, inserted error, or corrected error. The data for correct responses are averaged over positions. Error bars indicate 95% confidence intervals based on Fisher's least significant difference test calculated from the interaction between response type and error position. (**B**) Mean probabilities of reporting correct, error, inserted error, or corrected error for correct responses (correct), actual errors (error), inserted errors (inserted), and corrected errors (corrected).

task (correct, error, inserted error, and corrected error) to allow typists to distinguish sources of errors and correct responses and, therefore, provide a stronger test of illusions of authorship. We asked 24 skilled typists (WPM = 70.7 ± 16.4) to type 600 words, each of which was followed by a four-alternative explicit report screen. Typists typed 91.8% of the words correctly. Mean interkeystroke intervals, plotted in Fig. 3A, show post-error slowing for incorrect responses ($F_{1,138} = 117.7, p < 0.01$) and corrected errors ($F_{1,138} = 120.0, p < 0.01$), but not for inserted errors ($F < 1.0$), indicating that inner-loop detection distinguishes between actual errors and correct responses.

Explicit detection probabilities, plotted in Fig. 3B, show good discrimination between correct and error responses. For correct responses, typists said "correct" more than "error" [$t(23) = 97.29$, $p < 0.01$]; for error responses, typists said "error" more than "correct" [$t(23) = 8.22$, $p < 0.01$]. Typists distinguished actual errors from inserted errors well, avoiding an illusion of authorship. They said "error" more than "inserted" for actual errors [$t(23) = 7.06$, $p < 0.01$] and "inserted" more than "error" for inserted errors [$t(23) = 14.75$, $p < 0.01$]. However, typists showed a strong illusion of authorship with corrected errors. They were just as likely to call them correct responses as corrected errors [$t(23) = 1.38$].

The post-error slowing and post-trial report data show a dissociation between inner- and outer-loop error detection. We assessed the dissociation further by comparing post-error slowing on trials in which typists did and did not experience illusions of authorship (*21*). The pattern of post-error slowing was the same for both sets of trials (fig. S6), suggesting that the pattern in Fig. 3A is representative of all trials.

The three experiments found strong dissociations between explicit error reports and post-error slowing. These dissociations are consistent with the hierarchical error-detection mechanism that we proposed, with an outer loop that mediates explicit reports and an inner loop that mediates post-error slowing. This nested-loop description of error detection is consistent with hierarchical models of cognitive control in typewriting (*9*, *10*, *15–17*) and with models of hierarchical control in other complex tasks (2, *8*, *22*). Speaking, playing music, and navigating through space may all involve inner loops that take care of the details of performance (e.g., uttering phonemes, playing notes, and walking) and outer loops that ensure that intentions are fulfilled (e.g., messages communicated, songs performed, and destinations reached). Hierarchical control may be prevalent in highly skilled performers who have had enough practice to develop an autonomous inner loop. Previous studies of error detection in simple tasks may describe inner-loop processing. The novel contribution of our research is to dissociate the outer loop from the inner loop.

The three experiments demonstrate cognitive illusions of authorship in skilled typewriting (*11–14*). Typists readily take credit for correct output on the screen, interpreting corrected errors as their own correct responses. They take the blame for inserted errors, as in the first and second experiments, but they also blame the computer, as in the third experiment. These illusions are consistent with the hierarchical model of error detection, with the outer loop assigning credit and blame and the inner loop doing the work of typing (*10*, *17*). Thus, illusions of authorship may be a hallmark of hierarchical control systems (*2*, *11*, *22*, *23*).

References and Notes

1. P. M. A. Rabbitt, *J. Exp. Psychol.* **71**, 264 (1966).
2. D. A. Norman, *Psychol. Rev.* **88**, 1 (1981).
3. C. B. Holroyd, M. G. H. Coles, *Psychol. Rev.* **109**, 679 (2002).
4. N. Yeung, M. M. Botvinick, J. D. Cohen, *Psychol. Rev.* **111**, 931 (2004).
5. W. J. Gehring, B. Goss, M. G. H. Coles, D. E. Meyer, E. Donchin, *Psychol. Sci.* **4**, 385 (1993).
6. S. Dehaene, M. I. Posner, D. M. Tucker, *Psychol. Sci.* **5**, 303 (1994).
7. C. S. Carter *et al.*, *Science* **280**, 747 (1998).
8. K. S. Lashley, in *Cerebral Mechanisms in Behavior*, L. A. Jeffress, Ed. (Wiley, New York, 1951), pp. 112–136.
9. T. A. Salthouse, *Psychol. Bull.* **99**, 303 (1986).
10. G. D. Logan, M. J. C. Crump, *Psychol. Sci.* **20**, 1296 (2009).
11. T. I. Nielsen, *Scand. J. Psychol.* **4**, 225 (1963).
12. M. M. Botvinick, J. D. Cohen, *Nature* **391**, 756 (1998).
13. D. M. Wegner, *The Illusion of Conscious Will* (MIT Press, Cambridge, MA, 2002).
14. G. Knoblich, T. T. J. Kircher, *J. Exp. Psychol. Hum. Percept. Perform.* **30**, 657 (2004).
15. D. E. Rumelhart, D. A. Norman, *Cogn. Sci.* **6**, 1 (1982).
16. L. H. Shaffer, *Psychol. Rev.* **83**, 375 (1976).
17. X. Liu, M. J. C. Crump, G. D. Logan, *Mem. Cognit.* **38**, 474 (2010).
18. A. M. Gordon, J. F. Soechting, *Exp. Brain Res.* **107**, 281 (1995).
19. J. Long, *Ergonomics* **19**, 93 (1976).
20. P. Rabbitt, *Ergonomics* **21**, 945 (1978).
21. Materials and methods are available as supporting material on *Science* Online.
22. M. M. Botvinick, *Trends Cogn. Sci.* **12**, 201 (2008).
23. R. Cooper, T. Shallice, *Cogn. Neuropsychol.* **17**, 297 (2000).

24. We thank J. D. Schall for comments on the manuscript. This research was supported by grants BCS 0646588 and BCS 0957074 from the NSF.

Supporting Online Material
www.sciencemag.org/cgi/content/full/330/6004/683/DC1
Materials and Methods
SOM Text
Figs. S1 to S6
References

5 April 2010; accepted 13 September 2010
10.1126/science.1190483

Evidence for a Collective Intelligence Factor in the Performance of Human Groups

Anita Williams Woolley,[1]* Christopher F. Chabris,[2,3] Alex Pentland,[3,4] Nada Hashmi,[3,5] Thomas W. Malone[3,5]

Psychologists have repeatedly shown that a single statistical factor—often called "general intelligence"—emerges from the correlations among people's performance on a wide variety of cognitive tasks. But no one has systematically examined whether a similar kind of "collective intelligence" exists for groups of people. In two studies with 699 people, working in groups of two to five, we find converging evidence of a general collective intelligence factor that explains a group's performance on a wide variety of tasks. This "*c* factor" is not strongly correlated with the average or maximum individual intelligence of group members but is correlated with the average social sensitivity of group members, the equality in distribution of conversational turn-taking, and the proportion of females in the group.

As research, management, and many other kinds of tasks are increasingly accomplished by groups—working both face-to-face and virtually (*1–3*)—it is becoming ever more important to understand the determinants of group performance. Over the past century, psychologists made considerable progress in defining and systematically measuring intelligence in individuals (*4*). We have used the statistical approach they developed for individual intelligence to systematically measure the intelligence of groups. Even though social psychologists and others have studied for decades how well groups perform specific tasks (*5*, *6*), they have not attempted to measure group intelligence in the same way individual intelligence is measured—by assessing how well a single group can perform a wide range of different tasks and using that information to predict how that same group will perform other tasks in the future. The goal of the research reported here was to test the hypothesis that groups, like individuals, do have characteristic levels of intelligence, which can be measured and used to predict the groups' performance on a wide variety of tasks.

Although controversy has surrounded it, the concept of measurable human intelligence is based on a fact that is still as remarkable as it was to Spearman when he first documented it in 1904

[1]Carnegie Mellon University, Tepper School of Business, Pittsburgh, PA 15213, USA. [2]Union College, Schenectady, NY 12308, USA. [3]Massachusetts Institute of Technology (MIT) Center for Collective Intelligence, Cambridge, MA 02142, USA. [4]MIT Media Lab, Cambridge, MA 02139, USA. [5]MIT Sloan School of Management, Cambridge, MA 02142, USA.

*To whom correspondence should be addressed. E-mail: awoolley@cmu.edu